Game Theory

2.3 Iterated and Extensive Games

George Mason University, Spring 2018

To this point, we've been concerned with finding solutions to games that are played only once, but for some games those solutions do not hold when we allow for the game to be played more than once. We call these *iterated games*.

Any game can be iterated, of course, but we're going to examine only one famous example to bring out the ways in which iteration can change the game: Prisoner's Dilemma.

Ex. Prisoner's Dilemma (Raiffa, Flood and Dresher, Tucker).

Two individuals are know to have committed minor offenses but they are also suspected of robbing a bank together. They are arrested and placed in separate interrogation rooms. If they both confess to the robbery, then they will be sentenced to 5 years in prison. If only one confesses, then he will be set free but used as a character witness against the other individual who will receive a sentence of 10 years. If neither confesses, then they will both be convicted of the minor offenses and sentenced to only 1 year in prison.

	confess	do not confess
confess	-5,-5	0,-10
do not confess	-10,0	-1,-1

The only Nash equilibrium is $\langle \text{confess}, \text{confess} \rangle$.

Generalize. Of course, actually playing an iterated version of PD would take years...you've got to wait until both are released from prison and commit another crime. This isn't plausible.

But PD is just a name for a formal game that has a particular structure (notice the rows and columns are switched from the traditional presentation):

	cooperate	cover
cooperate	B,B	D,A
cover	A,D	0,0

Where $A > B > 0 > D$.
The only Nash equilibrium is $\langle \text{cover}, \text{cover} \rangle$, but the optimal strategy is $\langle \text{cooperate}, \text{cooperate} \rangle$

Many games that can be iterated have this structure. Consider, for example, a cartel setting production of a good.

Two ways of iterating a game:

Finitely iterated games are played a set number of times. *Each player knows when the next play of the game is the final play.*

Infinitely (or indefinitely) iterated games are not necessarily played an infinite number of times, *but neither player knows when the next play will be the last.*

Finitely Iterated PD.

Does the NE solution change when the game is indefinitely iterated?

Finitely Iterated PD.

Does the NE solution change when the game is indefinitely iterated?

No. If each player knows that the last play of the game is the last play of the game, then both will cover (or not cooperate) because they will see no chance of being rewarded for cooperation in future games and recognize that the other player is thinking the same thing.

Knowing that they will both cover in the final play, neither player will have reason to diverge from covering in the penultimate play. And so on...

Indefinitely Iterated PD.

The indefinite game is different. In this case, neither player knows when the next play is the last, and so both reason that there may be some benefit in cooperating given that the other player might *punish* you for covering in the next round of play.

If the other player has been relatively cooperative in past rounds of play, then one has reason to believe she will continue to be. But what if her cooperation is only a response to your cooperation? Testing that hypothesis might rock the boat and end up in a cycle of non-cooperation. Why risk those future gains by myopically playing a dominant strategy in this single round of play?

In fact, the long-term gains to be had from cooperation are potentially very large, so it might be a good strategy to venture cooperation on the first round of play to see if those gains can be had.

This is *tit-for-tat*.

Tit-for-Tat.

Always cooperate on the first round, and thereafter adjust your behavior to whatever your opponent did in the previous round.

Danger. If both players are playing tit-for-tat, then once someone defects, there's no chance for recovery because the other player will respond with non-cooperation which will in turn engender non-cooperation. What if defection was an accident or a momentary lapse? We've now lost all potential for future gains from cooperation.

Solution. A better strategy may be tit-for-tit-for-tat. This can come in two forms:

- **Pure.** Always cooperate on the first round, and thereafter, if your opponent covers, cooperate in the next round. If your opponent covers again, then cover in the next round.
- **Probabilistic.** Always cooperate on the first round, and thereafter respond to your opponents covering by cooperating with a probability of p and covering with a probability of $1 - p$.

Is tit-for-tat a rational strategy? We don't really know. It does well empirically, but there's no proof that it is a rational strategy for playing the game. It leads to a NE, but so does just about any other strategy...

Extensive Games.

Interactive choice situations often have a sequential structure. This structure is made explicit in *extensive games*.

An extensive game has *perfect information* if each player, when making his or her moves, knows everything that has previously occurred in the game.

An extensive game has *imperfect information* if each player, when making his or her moves, may have only partial information about what has previously occurred in the game.

We consider only games with perfect information here.

Ex. Sharing Apples.

There are two apples. I will propose an allocation of these apples to you and me. You will decide whether to accept or reject it. If you reject my proposed allocation, then neither of us gets any apples.

Ex. Chain-Store Game (Selten).

You own a chain-store with branches in fifty different U.S. cities. In each city, you face a single potential competitor who might encroach on your turf. If a competitor moves in, you can choose to fight or cooperate. Since fighting is costly, you prefer cooperating to fighting in the face of competition, but your favored outcome is when a potential rival does not move in at all. A potential rival is best off when they move in and you cooperate, but they would rather not compete than fight with you. Now suppose that this all plays out sequentially, city after city, and each potential competitor knows what has happened previously. What is your strategy?

Def 2.3.1. An *extensive game with perfect information* $\mathcal{G} = \langle \mathcal{N}, \mathcal{H}, \mathcal{P}, \mathcal{O}, g, \{\succcurlyeq_i\}_{i \in \mathcal{N}} \rangle$ consists of the following ingredients:

- a finite set of players $\mathcal{N}$
- a set of *histories* $\mathcal{H}$ (sequences of actions) where some subset of these $\mathcal{Z} \subseteq \mathcal{H}$ are the *terminal histories*
- a *player function* $\mathcal{P} : \mathcal{H} \setminus \mathcal{Z} \to \mathcal{N}$ that sends each non-terminal history $h \in \mathcal{H} \setminus \mathcal{Z}$ to the player $i \in \mathcal{N}$ who acts at this juncture
- a set of outcomes $\mathcal{O}$
- a function $g : \mathcal{Z} \to \mathcal{O}$ that sends each terminal history $h \in \mathcal{Z}$ to an outcome $o \in \mathcal{O}$
- a preference relation $\succcurlyeq_i$ for each player over the set of outcomes $\mathcal{O}$

A history $h = \langle a_1, a_2, ... \rangle$ is a sequence of actions by the players of the game (this can be infinite).

$\mathcal{A}(h) = \{a : \langle h, a \rangle \in \mathcal{H}\}$ is the set of actions available after h.

The set of histories $\mathcal{H}$ must satisfy the following three properties:

- $\emptyset \in \mathcal{H}$.
- Every initial sequence of a history is a history: if $\langle a_1, ..., a_n, ... \rangle \in \mathcal{H}$, then $\langle a_1, ..., a_n \rangle \in \mathcal{H}$.
- If every finite initial sequence of an infinite sequence is a history, then the infinite sequence is a history.

Def 2.3.2. A history $h \in \mathcal{H}$ is *terminal* just in case either h is infinite or $h = \langle a_1, ..., a_n \rangle$ and there is no a_{n+1} such that $\langle a_1, ..., a_n, a_{n+1} \rangle \in \mathcal{H}$.

A terminal history is a kind of sequential action profile for all players.

Def 2.3.3. $\mathcal{G}$ is *finite* iff $|\mathcal{H}|$ is finite.

Def 2.3.4. $\mathcal{G}$ has a *finite horizon* iff the longest history in $\mathcal{H}$ is finite.

The extensive game starts at $\emptyset$.

$\mathcal{P}(\emptyset)$ is the player who acts first.

$\mathcal{P}(\emptyset)$ chooses an action in $\mathcal{A}(\emptyset) = \{a : \langle \emptyset, a \rangle \in \mathcal{H}\}$.

Suppose that $\mathcal{P}(\emptyset)$ chooses $a_1 \in \mathcal{A}(\emptyset)$.

$\mathcal{P}(\langle \emptyset, a_1 \rangle)$ is the player who acts next.

$\mathcal{P}(\langle \emptyset, a_1 \rangle)$ chooses an action in $\mathcal{A}(\langle \emptyset, a_1 \rangle) = \{a : \langle \emptyset, a_1, a \rangle \in \mathcal{H}\}$.

Suppose that $\mathcal{P}(\langle \emptyset, a_1 \rangle)$ chooses $a_2 \in \mathcal{A}(\langle \emptyset, a_1 \rangle)$.

$\mathcal{P}(\langle \emptyset, a_1, a_2 \rangle)$ is the player who acts next.

And so on.

If $\mathcal{G}$ has a finite horizon, then this process eventually culminates in a terminal history $\langle \emptyset, a_1, ..., a_n \rangle$.

$g(\langle \emptyset, a_1, ..., a_n \rangle) \in \mathcal{O}$ is the outcome of the game.

The game $\mathcal{G}$ of Sharing Apples consists of:

- $\mathcal{N} = \{1, 2\}$
- $\mathcal{H} = \{\emptyset, \langle\emptyset, 20\rangle, \langle\emptyset, 11\rangle, \langle\emptyset, 02\rangle, \langle\emptyset, 20, a\rangle, \langle\emptyset, 20, r\rangle, \langle\emptyset, 11, a\rangle,$ $\langle\emptyset, 11, r\rangle, \langle\emptyset, 02, a\rangle, \langle\emptyset, 02, r\rangle\}$
  where
  $\mathcal{H} \setminus \mathcal{Z} = \{\emptyset, \langle\emptyset, 20\rangle, \langle\emptyset, 11\rangle, \langle\emptyset, 02\rangle\}$
  $\mathcal{Z} = \{\langle\emptyset, 20, a\rangle, \langle\emptyset, 20, r\rangle, \langle\emptyset, 11, a\rangle, \langle\emptyset, 11, r\rangle, \langle\emptyset, 02, a\rangle, \langle\emptyset, 02, r\rangle\}$
- $\mathcal{P}(\emptyset) = 1$, $\mathcal{P}(\langle\emptyset, 20\rangle) = \mathcal{P}(\langle\emptyset, 11\rangle) = \mathcal{P}(\langle\emptyset, 02\rangle) = 2$
- $\mathcal{O} = \{20, 02, 11, 00\}$
- $g(\langle\emptyset, 20, a\rangle) = 20$
  $g(\langle\emptyset, 20, r\rangle) = 00$
  $g(\langle\emptyset, 11, a\rangle) = 11$
  And so forth.
- $20 \succ_1 11 \succ_1 02 \sim_1 00$
  $02 \succ_2 11 \succ_2 20 \sim_2 00$

An extensive game can be visualized in a tree.

For example, the game $\mathcal{G}$ of Sharing Apples corresponds to the following *game tree*.

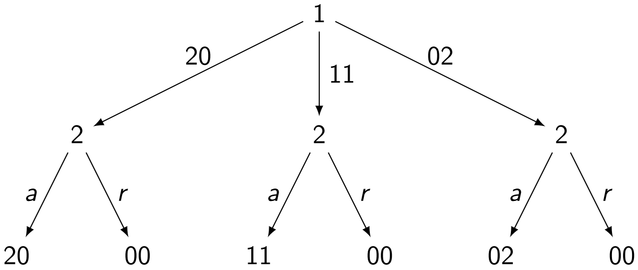

In an extensive game, players can have plans for how to act as the game unfolds.

Def 2.3.5. A *strategy* s_i of player $i \in \mathcal{N}$ is a function that maps each history $h \in \mathcal{H} \setminus \mathcal{Z}$ such that $\mathcal{P}(h) = i$ to an action $a \in \mathcal{A}(h)$.

That is, a strategy dictates how a player will act at *any* potential point in the game where this player is called to action.

In Sharing Apples, a strategy for player 1 will dictate how the apples are allocated. For example, $s_1(\langle\emptyset\rangle) = 11$.

A strategy for player 2 will dictate whether to accept or reject any proposed allocation. For example, $s_2(\langle\emptyset, 20\rangle) = r$, $s_2(\langle\emptyset, 11\rangle) = a$, and $s_2(\langle\emptyset, 02\rangle) = a$.

A strategy must specify an action at every choice point, even those that are unreachable if this very strategy is followed.

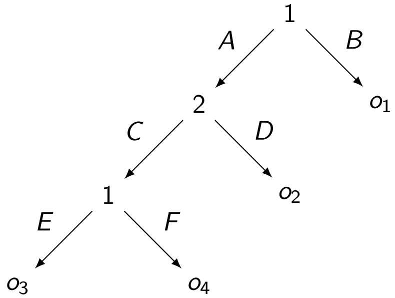

Player 1 has four strategies: AE, AF, BE, and BF (N.B. I will often use abbreviations for strategies).

BE and BF specify an action after history $\langle \emptyset, A, C \rangle$ even though this history will not be actualized if player 1 initially performs action B.

Let $\mathcal{S}_i$ designate the set of strategies for player $i \in \mathcal{N}$.

We will now work with strategy profiles in $\times_{i \in \mathcal{N}} \mathcal{S}_i$.

The outcome $o(s)$ of a strategy profile $s = \langle s_i, ..., s_{|\mathcal{N}|} \rangle$ is the outcome $o \in \mathcal{O}$ that results when each player $i \in \mathcal{N}$ follows their strategy s_i in this profile.

Formally, $o(s) = g(\langle a_1, a_2, .. \rangle)$ where $a_{n+1} = s_{\mathcal{P}(\langle a_1, ..., a_n \rangle)}(\langle a_1, ..., a_n \rangle)$.

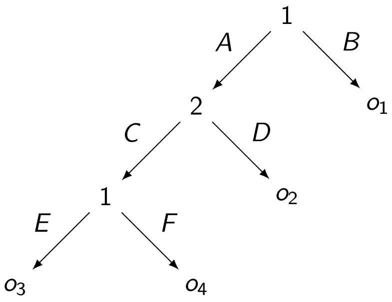

$$o(\langle AE, C\rangle) = ?$$

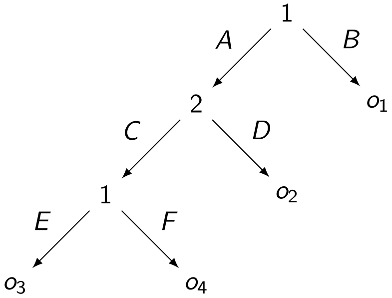

$$o(\langle AE, C\rangle) = o_3$$

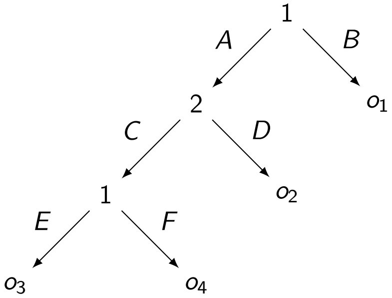

$$o(\langle AF, D\rangle) = ?$$

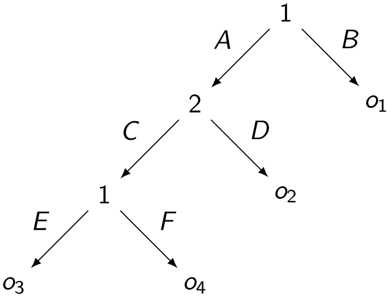

$$o(\langle AF, D\rangle) = o_2$$

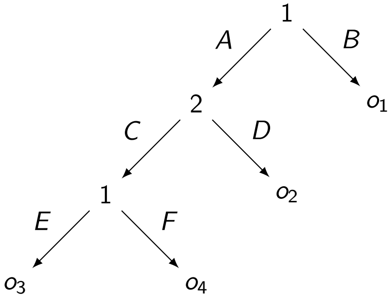

$$o(\langle BE, C\rangle) = ?$$

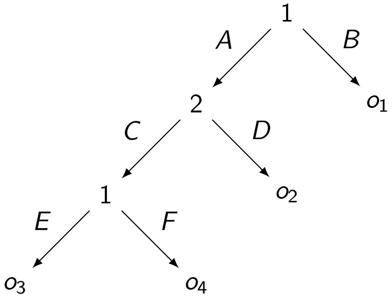

$$o(\langle BE, C \rangle) = o_1$$

Solution concepts such as *Nash equilibrium* for extensive games can be formulated in terms of strategy profiles.

Def 2.3.6. A Nash equilibrium of an extensive game with perfect information $\mathcal{G}$ is a strategy profile $s^* \in \times_{i \in \mathcal{N}} \mathcal{S}_i$ such that for every player $j \in \mathcal{N}$, the following condition holds:

$o(\langle s_j^*, s_{-j}^* \rangle) \succsim_j o(\langle s_j, s_{-j}^* \rangle)$ for each $s_j \in \mathcal{S}_j$.

In other words, given the other players' equilibrium strategy profile s_{-j}^*, the equilibrium strategy s_j^* of player j is optimal.

As before, we can also work with utilities.

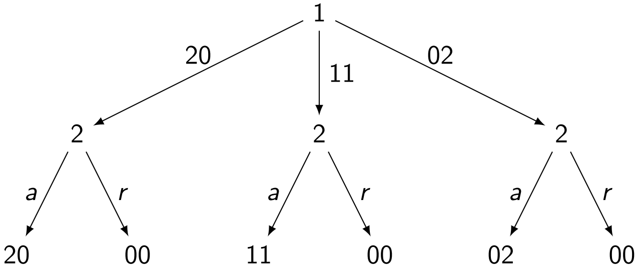

Assume that an apple is valued at a utile.

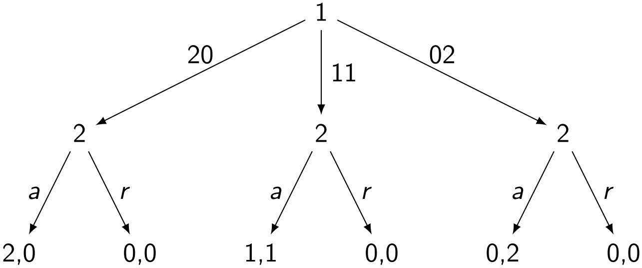

What are the Nash equilibria of Sharing Apples?

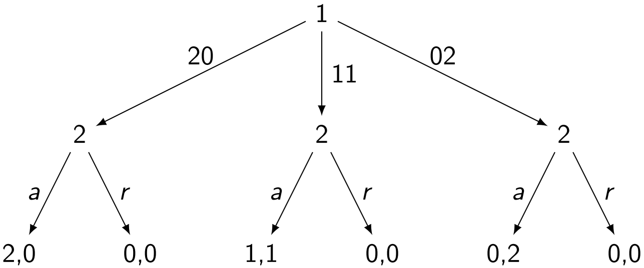

$$\langle 20, aaa\rangle, \langle 20, aar\rangle, \langle 20, ara\rangle, \langle 20, arr\rangle, \langle 20, rra\rangle, \langle 20, rrr\rangle,$$
$$\langle 11, raa\rangle, \langle 11, rar\rangle$$
$$\langle 02, rra\rangle$$

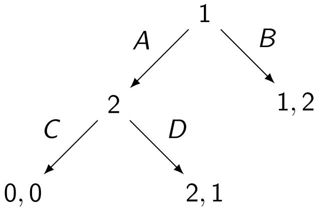

What are the Nash equilibria of this game?

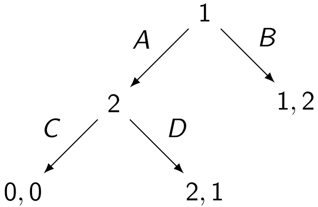

$$\langle A, D \rangle,\ \langle B, C \rangle$$

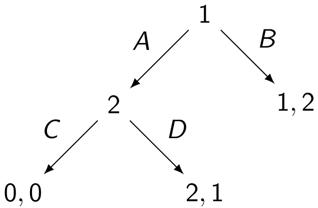

But $\langle B, C \rangle$ is not a stable solution to this game. It is sustained by the threat of player 2 to perform action C. But this action is irrational once player 1 has performed action A.

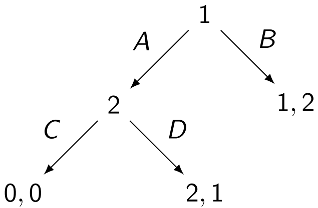

Moral. The solution concept of Nash equilibrium is ill-suited for extensive games since it ignores the sequential structure of these games.

Def 2.3.7. The *subgame* of the extensive game with perfect information $\mathcal{G} = \langle \mathcal{N}, \mathcal{H}, \mathcal{P}, \mathcal{O}, g, \{\succcurlyeq_i\}_{i \in \mathcal{N}} \rangle$ *that follows history h* is the game $\mathcal{G}(h) = \langle \mathcal{N}, \mathcal{H}|_h, \mathcal{P}|_h, \mathcal{O}, g|_h, \{\succcurlyeq_i\}_{i \in \mathcal{N}} \rangle$ consisting of the following ingredients:

- a finite set of players $\mathcal{N}$
- a set of histories $\mathcal{H}|_h$ where $h' \in \mathcal{H}|_h$ iff $\langle h, h' \rangle \in \mathcal{H}$
- a player function $\mathcal{P}|_h$ where $\mathcal{P}|_h(h') = \mathcal{P}(\langle h, h' \rangle)$
- a set of outcomes $\mathcal{O}$
- a function $g|_h : \mathcal{Z}|_h \to \mathcal{O}$ where $g|_h(h') = g(\langle h, h' \rangle)$
- a preference relation $\succcurlyeq_i$ for each player over the set of outcomes $\mathcal{O}$

The game $\mathcal{G}$ of Sharing Apples consists of:

- $\mathcal{N} = \{1, 2\}$
- $\mathcal{H} = \{\emptyset, \langle\emptyset, 20\rangle, \langle\emptyset, 11\rangle, \langle\emptyset, 02\rangle, \langle\emptyset, 20, a\rangle, \langle\emptyset, 20, r\rangle, \langle\emptyset, 11, a\rangle,$
  $\langle\emptyset, 11, r\rangle, \langle\emptyset, 02, a\rangle, \langle\emptyset, 02, r\rangle\}$
  where
  $\mathcal{H} \setminus \mathcal{Z} = \{\emptyset, \langle\emptyset, 20\rangle, \langle\emptyset, 11\rangle, \langle\emptyset, 02\rangle\}$
  $\mathcal{Z} = \{\langle\emptyset, 20, a\rangle, \langle\emptyset, 20, r\rangle, \langle\emptyset, 11, a\rangle, \langle\emptyset, 11, r\rangle, \langle\emptyset, 02, a\rangle, \langle\emptyset, 02, r\rangle\}$
- $\mathcal{P}(\emptyset) = 1$, $\mathcal{P}(\langle\emptyset, 20\rangle) = \mathcal{P}(\langle\emptyset, 11\rangle) = \mathcal{P}(\langle\emptyset, 02\rangle) = 2$
- $\mathcal{O} = \{20, 02, 11, 00\}$
- $g(\langle\emptyset, 20, a\rangle) = 20$
  $g(\langle\emptyset, 20, r\rangle) = 00$
  $g(\langle\emptyset, 11, a\rangle) = 11$
  And so forth.
- $20 \succ_1 11 \succ_1 02 \sim_1 00$
  $02 \succ_2 11 \succ_2 20 \sim_1 00$

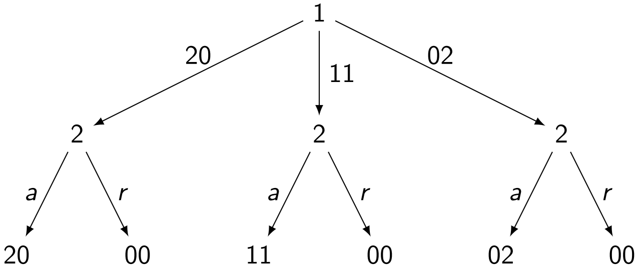
1
20
11
02
2
2
2
a
r
a
r
a
r
20
00
11
00
02
00

The subgame $\mathcal{G}(\langle\emptyset, 11\rangle)$ of Sharing Apples that follows history $\langle\emptyset, 11\rangle$ consists of:

- $\mathcal{N} = \{1, 2\}$
- $\mathcal{H} = \{\emptyset, \langle\emptyset, a\rangle, \langle\emptyset, r\rangle\}$
  where
  $\mathcal{H} \setminus \mathcal{Z} = \{\emptyset\}$
  $\mathcal{Z} = \{\langle\emptyset, a\rangle, \langle\emptyset, r\rangle\}$
- $\mathcal{P}(\emptyset) = 2$
- $\mathcal{O} = \{20, 02, 11, 00\}$
- $g(\langle\emptyset, a\rangle) = 11$
  $g(\langle\emptyset, r\rangle) = 00$
- $20 \succ_1 11 \succ_1 02 \sim_1 00$
  $02 \succ_2 11 \succ_2 20 \sim_2 00$

$$
\begin{tikzcd}
& 2 \arrow[dl, "a"'] \arrow[dr, "r"] & \\
11 & & 00
\end{tikzcd}
$$

Given a strategy s_i for player $i \in \mathcal{N}$ and a history $h \in \mathcal{H}$ in the extensive game $\mathcal{G}$, let $s_i|_h$ designate the strategy induced by s_i in the subgame $\mathcal{G}(h)$. That is, $s_i|_h(h') = s_i(\langle h, h' \rangle)$ for each $h' \in \mathcal{H}|_h$.

Def 2.3.8. The *subgame perfect equilibrium* of an extensive game with perfect information $\mathcal{G} = \langle \mathcal{N}, \mathcal{H}, \mathcal{P}, \mathcal{O}, g, \{\succcurlyeq_i\}_{i \in \mathcal{N}} \rangle$ is a strategy profile $s^* \in \times_{i \in \mathcal{N}} \mathcal{S}_i$ such that for every player $j \in \mathcal{N}$ and non-terminal history $h \in \mathcal{H} \setminus \mathcal{Z}$ for which $\mathcal{P}(h) = j$, the following condition holds:

$o(\langle s_j^*|_h, s_{-j}^*|_h \rangle) \succcurlyeq_j o(\langle s_j|_h, s_{-j}^*|_h \rangle)$ for each $s_j \in \mathcal{S}_j$.

Equivalently, s^* is a subgame perfect equilibrium of $\mathcal{G}$ just in case $s^*|_h$ is a Nash equilibrium of $\mathcal{G}(h)$ for each $h \in \mathcal{H} \setminus \mathcal{Z}$.

A subgame perfect equilibrium is a Nash equilibrium but the converse needn't be true.

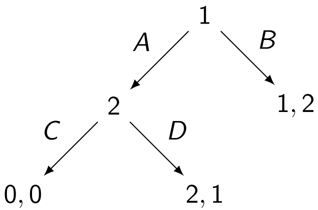

$\langle A, D\rangle$ and $\langle B, C\rangle$ are Nash equilibria.

Are either of these subgame perfect equilibria?

Only $\langle A, D\rangle$ since D but not C is a Nash equilibrium of $\mathcal{G}(\langle\emptyset, A\rangle)$.

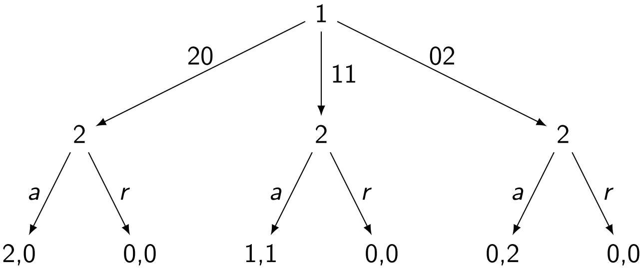

$\langle 20, aaa\rangle$,$\langle 20, aar\rangle$,$\langle 20, ara\rangle$,$\langle 20, arr\rangle$,$\langle 20, rra\rangle$,$\langle 20, rrr\rangle$,
$\langle 11, raa\rangle$,$\langle 11, rar\rangle$,$\langle 02, rra\rangle$ are Nash equilibria.

Are any of these subgame perfect equilibria?

Only $\langle 20, aaa\rangle$ and $\langle 11, raa\rangle$.

Thm 2.3.1. Every finite extensive game with perfect information has a subgame perfect equilibrium.

Since subgame perfect equilibria are Nash equilibria, every finite extensive game with perfect information has a Nash equilibrium.

Example: chess with the rule that the game is a draw when a board position is repeated three times.

Ex. Chain-Store Game (Selten).

You own a chain-store with branches in fifty different U.S. cities. In each city, you face a single potential competitor who might encroach on your turf. If a competitor moves in, you can choose to fight or cooperate. Since fighting is costly, you prefer cooperating to fighting in the face of competition, but your favored outcome is when a potential rival does not move in at all. A potential rival is best off when they move in and you cooperate, but they would rather not compete than fight with you. Now suppose that this all plays out sequentially, city after city, and each potential competitor knows what has happened previously.

What are the subgame perfect equilibria of this game?

The subgame tree for any particular city (where you are player 1 and your rival is player 2):

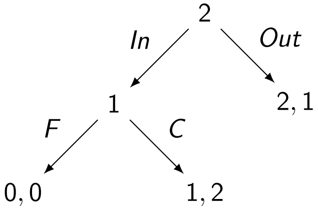

The only subgame perfect equilibrium is $\langle In, C \rangle$.

In the full Chain-Store Game, the unique subgame perfect equilibrium is where the rival enters each city and you cooperate.

But isn't it better to fight at first and develop a reputation for being aggressive?

Ex. Centipede Game (Rosenthal).

$$\begin{tikzcd}
1 \arrow[r] \arrow[d] & 2 \arrow[r] \arrow[d] & 1 \arrow[r] \arrow[d] & 2 \arrow[r] \arrow[d] & 1 \arrow[r] \arrow[d] & 2 \arrow[r] \arrow[d] & 6,5 \\
1,0 & 0,2 & 3,1 & 2,4 & 5,3 & 4,6 &
\end{tikzcd}$$

What are the subgame perfect equilibria of this game?

Each player stops the game and cashes out in each period.

Moral. It is useful to develop incentive mechanisms for cooperation.